

# Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-06 06:12:16

PAGE 1

REFERENCE NO: 311

This contribution was submitted to the National Science Foundation as part of the NSF CI 2030 planning activity through an NSF Request for Information, [https://www.nsf.gov/publications/pub\\_summ.jsp?ods\\_key=nsf17031](https://www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf17031). Consideration of this contribution in NSF's planning process and any NSF-provided public accessibility of this document does not constitute approval of the content by NSF or the US Government. The opinions and views expressed herein are those of the author(s) and do not necessarily reflect those of the NSF or the US Government. The content of this submission is protected by the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (<https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>).

## Author Names & Affiliations

- Alexander Tchekhovskoy - UC Berkeley

## Contact Email Address (for NSF use only)

(Hidden)

## Research Domain, discipline, and sub-discipline

Astrophysics Theory; Magnetized Fluid Dynamics; GPU-enabled Massively Parallel Magnetized Fluid Dynamics Simulations

## Title of Submission

Massively Parallel GPU Computing Addressing Extreme Scale Separation in Astrophysical Systems

## Abstract (maximum ~200 words).

By spanning more than ten orders of magnitude in distance, astrophysical systems present us with an enormous contrast in length scales that is an extreme challenge for numerical simulations. Observations of galaxies tell us that only a small fraction of the gas reservoir at the outskirts of galaxies reaches the central supermassive black holes. Does the gas experience a gravitational instability and form stars along the way? Or does it get turned around and expelled as an outflow instead of falling into the black hole? The answers to these questions are not agreed upon, and the computational resources available to us as a community are presently insufficient to address these questions. However, a confluence of new generation computational tools and supercomputers holds great promise at addressing these problems. In particular, graphical processing units (GPUs) show great promise at accelerating the computations (by more than an order of magnitude compared to CPUs) and offer a cost-effective way of powering next-generation supercomputing resources adequate for solving these challenging problems.

**Question 1** Research Challenge(s) (maximum ~1200 words): Describe current or emerging science or engineering research challenge(s), providing context in terms of recent research activities and standing questions in the field.

The vast discrepancy in scales intrinsically present in astronomical systems is astounding. In order to reach the supermassive black holes at the centers of galaxies, the ambient gas needs to traverse over 8 orders in magnitude in distance from the galaxy outskirts to the center. The central black holes, which consume the gas, produce collimated streams of gas and magnetic fields that carry the energy from the black holes and deposit it into the ambient medium 10 orders of magnitude in distance away. This energy deposition might heat up and expel the ambient gas and prevent it from feeding the black hole in the first place. These complex cycles of black hole growth and starvation might regulate the growth of the galaxy as a whole and shape the universe we live in. Despite being clearly important, this problem is extremely difficult to solve. This is because the length and time scales near the galaxy outskirts are tens of orders of magnitude

# Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-06 06:12:16

PAGE 2

REFERENCE NO: 311

longer than near the black hole, making it impossible to simulate such systems as a whole from first principles. A modern approach is to construct multi-scale models that treat the black hole physics of the central parts of the galaxy as a rather crude sub-grid model. The prohibitive scale separation makes it presently impossible to properly calibrate such sub-grid models of black hole feeding and feedback. However, next-generation supercomputers will enable direct attacks on such problems.

**Question 2** Cyberinfrastructure Needed to Address the Research Challenge(s) (maximum ~1200 words): Describe any limitations or absence of existing cyberinfrastructure, and/or specific technical advancements in cyberinfrastructure (e.g. advanced computing, data infrastructure, software infrastructure, applications, networking, cybersecurity), that must be addressed to accomplish the identified research challenge(s).

Recently, my team made a breakthrough in utilizing GPUs in order to simulate magnetized fluid dynamics around black holes. We showed that a modern GPU is capable of outperforming a modern CPU by more than an order of magnitude, bringing about a revolution in the study of black hole accretion. In fact, this is the first time such a substantial speedup is possible for fluid modeling on a curved space-time of a spinning black hole. Importantly, the simplicity of GPU architecture -- a collection of tens of thousands of threads each of which takes on a simple task -- is perfect for accelerating grid-based fluid codes such as ours. Without any changes to our code, its speed on a single GPU has doubled every 2 years for the past 4 years: the speed of our code on NVidia P100 GPUs (released in 2016) is 4 times that on NVidia K20x GPUs (released in 2012). From the specifications, we anticipate this speed-up trend to continue with the new generation of Volta GPUs whose release is slated for 2018. Importantly, our code shows very good parallel scaling to thousands of GPUs, primarily due to the locality of the included physics. This allows us to attack the most difficult and challenging problems, such as black hole feeding and feedback. With an eye toward the next decade, the ever increasing power of GPUs will make it possible to mount direct attacks on the problem of black hole feeding and feedback. Present NSF-funded resources are limited to the Blue Waters supercomputer, which has been our workhorse for pushing the studies of astrophysical systems to its extreme. However, with Blue Waters slated for retirement a little more than a year from now, there is a possibility of a vacuum in supercomputing resources. In the absence of a timely successor to Blue Waters, the progress on many high-impact projects would be stalled.

**Question 3** Other considerations (maximum ~1200 words, optional): Any other relevant aspects, such as organization, process, learning and workforce development, access, and sustainability, that need to be addressed; or any other issues that NSF should consider.

For many problems of interest codes need to be able to handle non-local effects, such as self-gravity, which requires faster interconnect between the GPUs in order to ensure satisfactory parallel performance. Thus, investment into the development of next-generation interconnect and its acquisition is crucial in enabling the sustained progress of massively parallel computing.

Most importantly, the availability of a community accessible resource, such as the successor(s) Blue Waters, is crucial for supporting the cutting edge computational science that pushes the limits of possibilities. A long-term plan for maintaining and upgrading such a resource would ensure sustained progress in addressing cutting edge scientific problems.

In order to ensure that the community is able to use the next-generation supercomputing resources, before locking the funding to a particular hardware architecture, it is important to form focus groups of key users and explore the challenges of porting the codes to new architectures and efficiently using these architectures. Establishing an online resource that would summarize the best practices and successes and failures of code conversion to different architectures across different fields of science would prove an invaluable resource for informed decision-making. This will ensure that the large investments of NSF into the cyberinfrastructure pay off and maximally benefit the scientific community. The present structure of the Blue Waters supercomputer, with separate CPU- and GPU-based partitions might be a model of a successful hardware architecture compromise that allows diverse community of researchers take advantage of the system.

## Consent Statement

- "I hereby agree to give the National Science Foundation (NSF) the right to use this information for the purposes stated above and to display it on a publically available website, consistent with the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0

# Submission in Response to NSF CI 2030 Request for Information

**DATE AND TIME:** 2017-04-06 06:12:16

**REFERENCE NO:** 311

**PAGE 3**

---

International License (<https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>)."

---